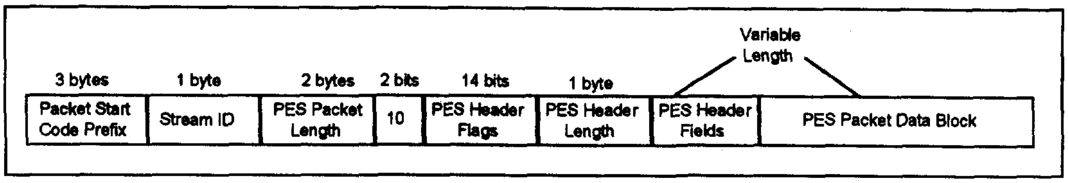

Figure 8.9. Structural overview of a PES packet.

Each elementary stream is identified by a unique stream_id. The PES packets formed from elementary streams supplied by each encoder carry the corresponding stream_id. PES packets carrying various types of elementary streams can be multiplexed to form a program or transport stream in accordance with the MPEG-2 Systems standard.

PES packets for video, including new PTS and DTS values, occur once every picture (or video access unit). The PES packets are also aligned to the first occurrence of a sequence, a GOP or a picture start code after the end of a picture, i.e., the first bytes seen within the payload of a video PES packet belong to either a sequence, a GOP or a picture start code. Further, new PES packet data always starts a new transport packet, and stuffing bytes are used in the adaptation header of the transport packets to ensure that PES packets always end on transport packet boundaries.

8.5.1 PES header flags

A breakdown of the PES header flags is shown in Figure 8.10. These flags are a combination of indicators of the properties of the bit stream and indicators of the existence of additional fields in the PES header. The following table describes the flags present in the header. The flags not supported by the system are set to '0' and form the basis of some of the "constraints" discussed earlier. (These entries are shaded in the table.)

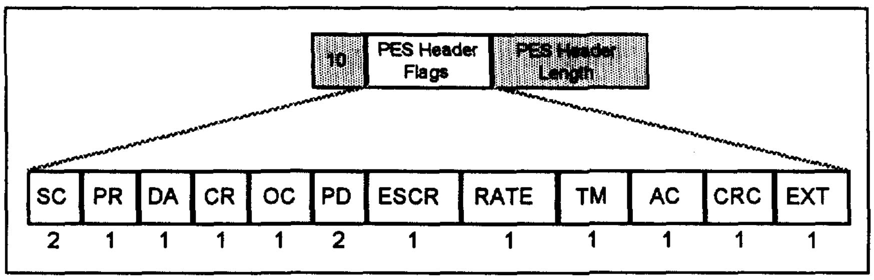

Figure 8.10. PES header flags in their relative positions (all sizes in bits).

8.5.2 The PES header

The PES header immediately follows the PES_header_length field, which indicates the header size in bytes. The size of the header includes all the header fields, any extension fields, and stuffing_bytes. The flags described in the previous Section indicate the organization of the PES header, i.e., which fields it does and does not contain. In essence, all the fields of the PES header are optional. Certain applications require particular fields to be set appropriately. For example, transport of video PES packets in the digital television system requires that the data_alignment_indicator be set. The trick mode flag is not set in this case. For DSM retrieval of video, the opposite is true. It is the application encoder's function to provide the information required to set the appropriate flags, and encode the corresponding fields. The fields are further described in the following Sections. The association between the flags and the corresponding fields is obvious.

The PES header fields are organized according to Figure 8.11 for the PES packets for video elementary streams. Most fields require marker bits to be inserted, as described later, in order to avoid the occurrence of long strings of 0's which could resemble a start code.

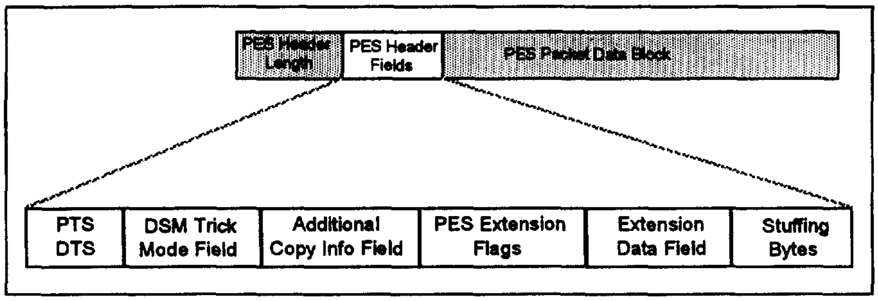

Figure 8.11. Organization of PES header.

8.5.2.1 PTS and DTS

The presentation_time_stamp (PTS) informs the decoder of the intended time of presentation of a presentation unit, and the decoding_time_stamp (DTS) is the intended time of decoding of an access unit. An access unit is an encoded presentation unit. When it is encoded, the PTS refers to the presentation unit corresponding to the first access unit occurring in the packet. If an access unit does not occur in a PES packet, the header shall not contain a PTS. An audio access unit occurs if the first byte of the synchronization word of an audio frame is present. A video access unit may be either a video sequence, a GOP, or a picture header as defined in Section 2.1.1 of ISO/IEC 13818-1. For terrestrial broadcast, each video frame is PES packetized. Under normal conditions, the DTS may be derived from the PTS. The DTS is not encoded when its value is equal to the value of the PTS, such as when the frame is a B-frame or when there are no B-frames in the sequence and hence no frame reordering delay. Under no circumstance does the DTS occur by itself; it must occur along with the PTS, although the converse is not true. The PTS field is organized as shown in Figure 8.12, if it present without the DTS.

If both the PTS and DTS are sent, the organization of Figure 8.13 is required. Here, the DTS field is defined in the same manner as the PTS field.

8.5.2.2 PES extension flags

The header may contain additional flags if the EXT flag bit (shown in Figure 8.10) is set. These flags are transmitted in a one byte data field as shown in Figure 8.14.

The flags indicate whether further extensions to the PES header exist. As with the flags defined previously, the flag is set to '1' if the header field is present.

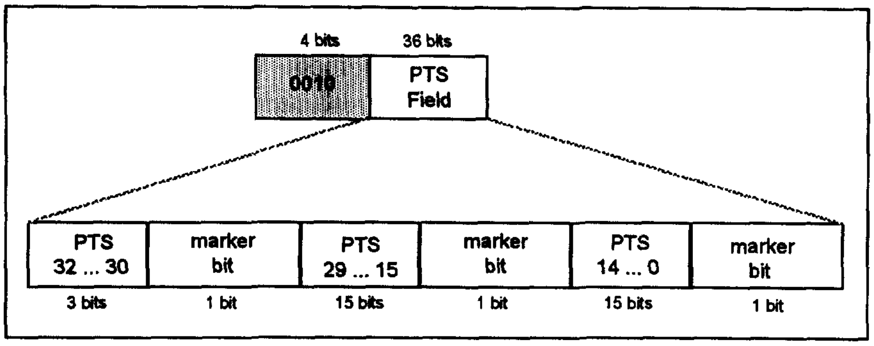

Figure 8.12. Organization of the PTS field when only the PTS is encoded.

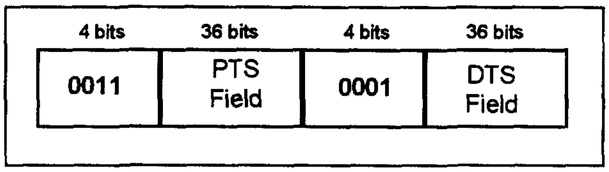

Figure 8.13. Organization of the PTS and DTS field when both PTS and DTS are encoded.

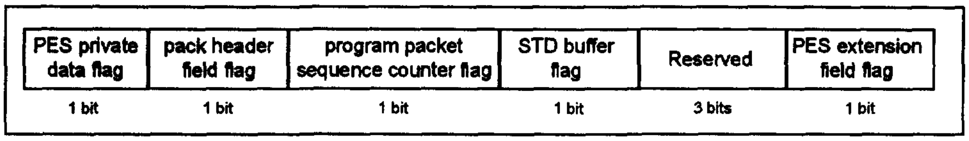

Figure 8.14. Organization of the PES extension flags field.

8.5.3 Conditional access

The transport protocol implements functions useful for supporting conditional access. The functionality that is available is flexible and complete in the sense of supporting all transmission aspects of applicable key encryption and scrambling approaches that may be used. Conditional access is also flexible in the sense that it can be applied separately to each elementary stream providing the ability to selectively scramble the different elementary streams (e.g., audio, video, etc.) in a program if desired.

A conditional access system operates on the principle of randomizing the transmitted data so that unauthorized decoders cannot decode the signal. Authorized decoders receive a "key" which initializes the circuit which inverts the bit randomization. In subsequent discussion, we use the term scrambling to mean the pseudo-random inversion of data bits based on a "key" which is valid for a short time. We use the term encryption to mean the process of transforming the "key" into an encrypted key by a means which protects the key from unauthorized users. From a cryptographic point of the view, this transformation of the key is the only part of the system which protects the data from a highly motivated pirate. The scrambling portion of the process alone, in the absence of key encryption, can be defeated. Conditional Access (CA) is a blanket term for the system which implements the key encryption and distribution. By virtue of the function of a CA system, the details remain proprietary, however the EIA and NCTA are working to standardize interfaces within consumer equipment to meet the needs of secure applications. The emerging interface standard is referred to as the National Renewable Security System (NRSS).

There are three features of the digital television transport system that support conditional access. The first feature is the two bit transport_scrambling_control field which signals the decoder whether the transport packet is scrambled or not. In the case that it is scrambled, the field identifies which of the scrambling keys was used. The second feature is the ability to insert "private" data at several places in the transport stream. These include entirely private streams and private fields in the adaptation header of the transport bit stream being scrambled. These private fields can be used to transmit the encrypted scrambling key to the decoding device. Thirdly, PID1 is used to identify those private streams which carry program authorizations.

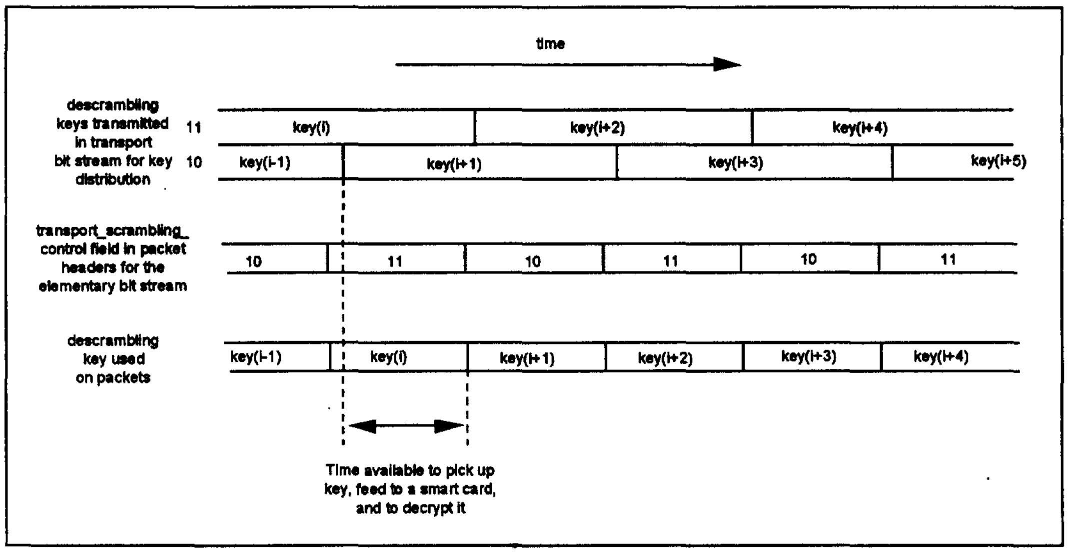

Figure 8.15. Illustration of key distribution and usage process.

An example of the key distribution process is shown in Figure 8.15. When the bit stream is scrambled, one descrambling key needs to be in use while the other is being received and decrypted. Two keys are transmitted at any time, with the keys being linked to a transport_scrambling_control value as shown in Figure 8.15. The transmission of a key should begin well before it is going to be used, to allow time to decrypt it. Note that this function does not bound the total number of keys that may be used during an entire transmission session. The proper key to be used to descramble is signaled in the transport prefix in the transport_scrambling_control field. The transport_scrambling_control takes on one of the 4 states shown in Table 8.1.

8.5.3.1 Conditional access example

Figure 8.16 illustrates an example of a digital television receiver that is operating in an environment which uses an NRSS device to implement the conditional access functions. The NRSS interface provides for 50 Mbps of throughput through a security device. The host microprocessor in the digital television receiver decoder communicates with the security device as to the desired program, etc., over the Serial In/Out control channel. The NRSS device itself filters the transport data stream, presented at the Din input, for all information which pertains to the CA functions. The secure processor decodes the CA information required to descramble the selected service. The security system algorithms and circuits decode the Entitlement Control Message (ECM) and Entitlement Management Message (EMM) information and supply the scrambling keys to the descrambling block, shown here as a Data Encryption Standard (DES) descrambler. The descrambled data is returned to the digital television receiver on the Dout output connector.

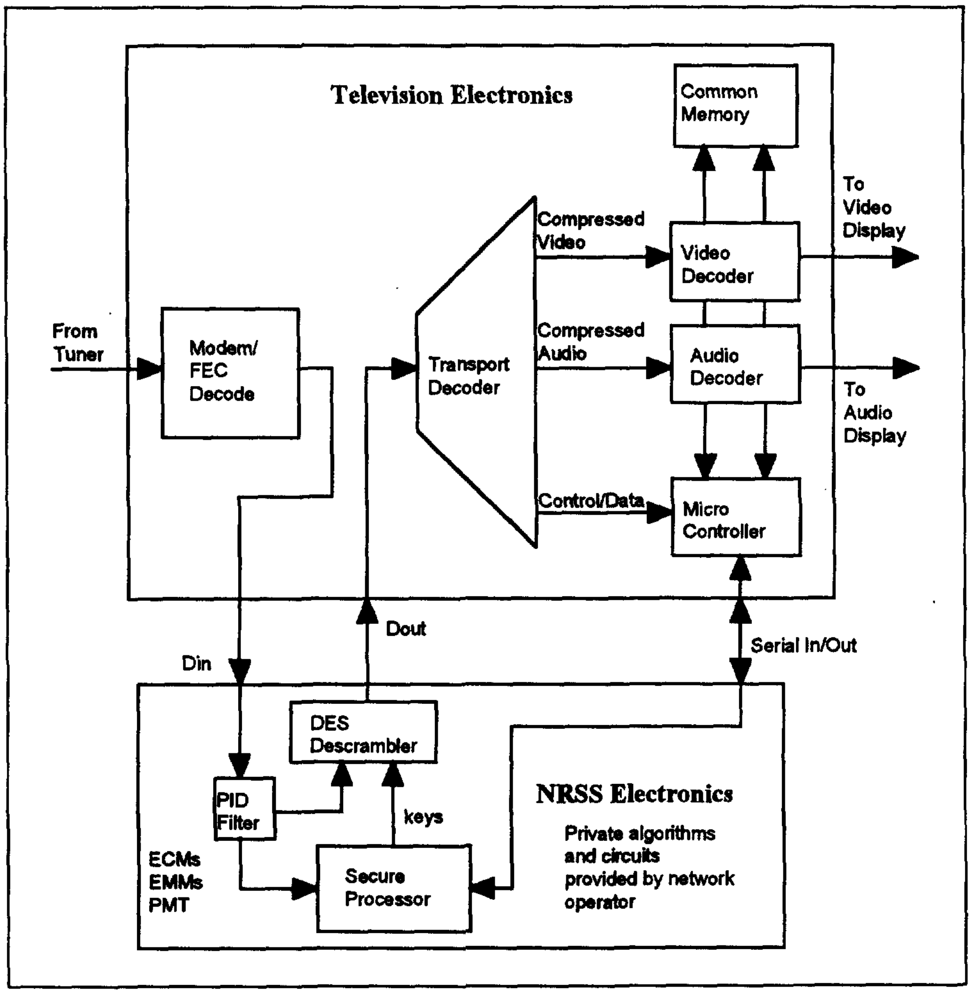

Figure 8.16. Example NRSS based A/V decoder in a secure service environment.

8.5.4 Compatibility with other transport systems

The transport system interoperates with two of the most important alternative transport systems. It is syntactically identical with the MPEG-2 transport stream definition, with the Digital Television Standard being a subset of the MPEG-2 specification. Annex C of the Digital Television Standard completely specifies the special constraints and requirements of the subset. The transport system also has a high degree of interoperability with the ATM definition being finalized for Broadband ISDN. Furthermore, as several of the cable television and Direct Broadcast Satellite (DBS) systems currently in use or being designed employ MPEG-2 transport layer syntax, the degree of interoperability with such deployed systems should be quite high (possibly requiring a translation if the cable television or DBS system deploys a slightly incompatible MPEG-2 variant).

Interoperability has two aspects. The first is syntactic and refers only to the coded representation of the digital television information. The second relates to the delivery of the bit stream in real time. This aspect of interoperability is beyond the scope of this discussion, but it should be noted that to guarantee interoperability with a digital television receiver conforming to the Standard, the output bit stream of the alternative transport system must have the proper real-time characteristics.

8.5.4.1 Interoperability with MPEG-2

In the development of the digital television transport specification, the intent has never been to limit the design by the scope of the MPEG-2 Systems definition. The system is interoperable with MPEG-2 decoders as the transport is a constrained subset of the MPEG-2 transport syntax. The constraints are imposed for reasons of increased performance with respect to channel acquisition, bandwidth efficiency and decoder complexity.

The system also supports bit streams and services beyond the compressed video and audio services, e.g., text-based services, emergency messages, and other future ancillary services. A means of identifying such bit streams is necessary, but is not part of the MPEG-2 definition. There is a method of encoding such a registration descriptor when an authority to administrate registration and catalog registered codes is identified. The method of encoding the registered value is by means of the registration_descriptor in the PSI stream.

8.5.4.2 Interoperability with ATM

The MPEG-2 transport packet size is such that it can be easily partitioned for transfer in a link layer that supports Asynchronous Transfer Mode (ATM) transmission. The MPEG-2 transport layer and the ATM layer serve different functions in a video delivery application. The MPEG-2 transport layer solves MPEG-2 presentation problems and performs the multi-media multiplexing function. The ATM layer solves switching and network adaptation problems.

There are several possible methods for mapping the MPEG-2 transport packet into the ATM format, and international standards organizations are standardizing the method to be used in different application domains. A popular method for constant bit rate (CBR) sources in the video-on-demand (VOD) application, under standardization by the ITU-T and the ATM Forum, is to perform cell aligned packet mapping with an AAL-5 Packet Data Unit (PDU). This technique is presented here to serve as an example.

8.5.4.2.1 ATM cell and transport packet structures

The ATM cell consists of two parts: a five-byte header and a forty-eight-byte information field. The header, primarily significant for networking purposes, consists of the fields shown in Table 8.4.

Table 8.4 ATM Cell Header Fields

GFC	A four bit Generic Flow Control field used to control the flow of traffic across the User Network Interface (UNI). Exact mechanisms for flow control are under investigation.
VPI	An eight bit network Virtual Path Identifier.
VCI	A sixteen bit network Virtual Circuit Identifier.
PT	A three bit Payload Type (i.e., user information type ID).
CLP	A one bit Cell Loss Priority flag (eligibility of the cell for discard by the network under congested conditions).
HEC	An eight bit Header Error Control field for ATM header error correction
AAL	ATM Adaptation Layer bytes (user specific header).

The ATM User Data Field consists of forty-eight bytes, where up to four of these bytes can be allocated to an Adaptation Layer.

Figure 8.17 illustrates the differences between the format of an ATM cell and the format of the MPEG-2 transport packet.

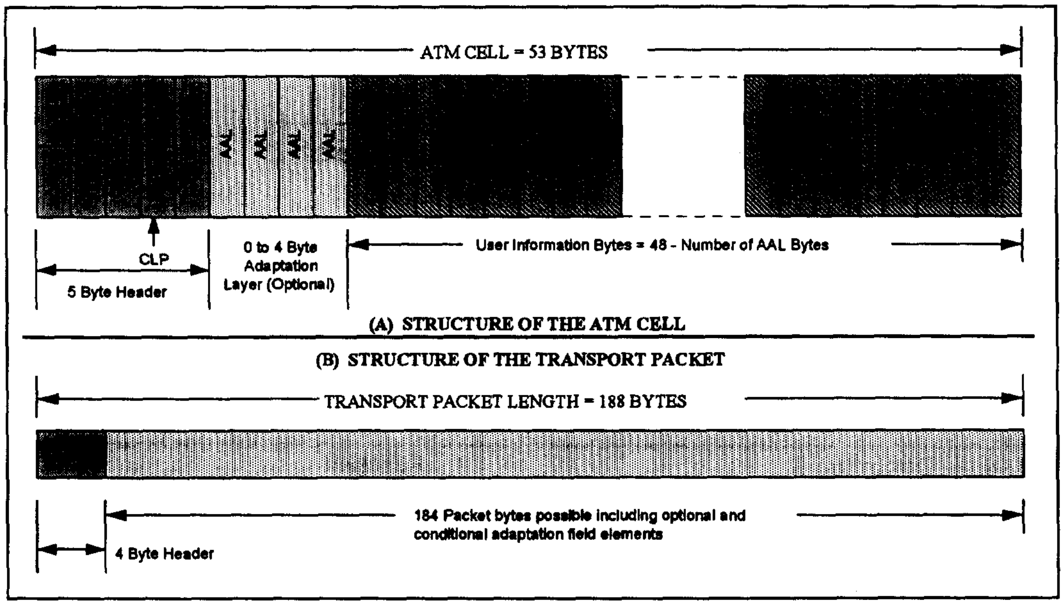

Figure 8.17. Comparison of the ATM cell structure and the MPEG-2 transport packet structure.

8.5.4.2.2 Mapping transport packets to an AAL-5 packet data unit (PDU)

A preferred solution to cell/packet alignment is to encapsulate one or two transport packets into a AAL-5 PDU structure. This is illustrated in Figure 8.18. Two

MPEG-2 transport packets, consisting of 376 bytes, are concatenated and the AAL-5 PDU trailer is computed and appended. The PDU is then segmented and the segments are inserted as the payload of 8 successive ATM cells. The ATM cell which contains the PDU header is identified by a special value in the PT field of the ATM cell header. Two processes can contribute to the time jitter of each MPEG-2 packet as it traverses the ATM network. The first is that ATM cells may be delayed at network nodes as part of the ATM cell scheduling and routing algorithms. This component is not controllable by the encoder-side equipment. The second component is the packetization delay in forming the 2-packet (8-cell) PDU. As the PCR value represents a real-time sample of the Program Clock Reference, jitter in packets containing the PCR is undesirable. The value of jitter is manifest in the decoder as a noise source at the input to the PLL which is used to reconstruct the audio and video sample clocks. This jitter can result in audible and visible artifacts if the effects are not properly mitigated. To ease the problem of jitter mitigation, it is common practice to add the following restriction to PDU creation: transport packets which carry a PCR sample shall be the last packet in a PDU. Consequently, if a new PDU is being formed and a PCR is in the first transport packet, then the encoder will issue a short (4 cell) PDU in order to limit unnecessary contributions to the PCR jitter.

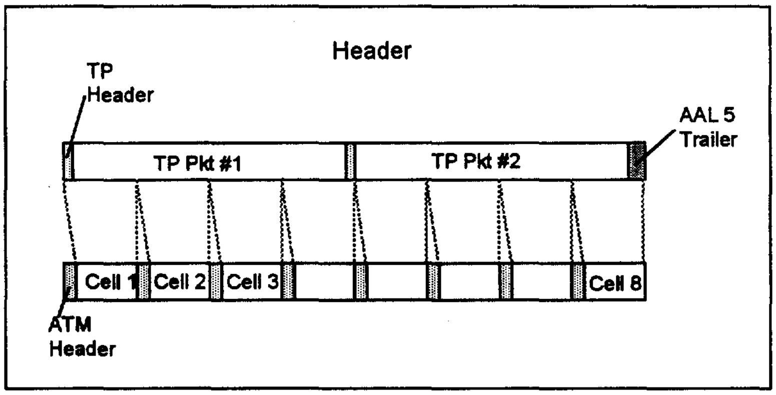

Figure 8.18. Mapping of two MPEG-2 transport packets into 8 ATM cells using the AAL-5 PDU.

Table 8.5 Elements of the AAL-5 PDU Trailer

PDU Field	Size(bytes)	Function
UU	1	User to User information to allow the encoder application to signal the decoder application with application dependent information.
CPI	1	Reserved for future use.
Length	2	16-bit binary number which represents the length of the PDU payload. It allows padding to follow the PDU payload in applications which require it.
CRC-32	4	Cyclic Redundancy Check calculated over the payload of the PDU.

The AAL 5 PDU trailer performs error detection and segmentation and reassembly (SAR) functions. The PDU trailer is an 8-byte structure, with the elements summarized in Table 8.5. The UU field may be defined by the ITU-T and/or the ATM Forum with specific information for the VOD application, as could the CPI field. This may be a long term enhancement, but early adopters of MPEG/ATM systems will probably restrict consideration to systems which use generic AAL-5 functions. The length field is important to locate any padding which may be required to pad the PDU to be an incremental number of ATM cells. Padding will be required when it is necessary to issue a short PDU because a PCR packet has arrived at the PDU assembly processor. The CRC-32 provides an end-to-end error detection function.

9. RF/TRANSMISSION SYSTEMS CHARACTERISTICS

9.1 Introduction and system overview

The VSB system has two modes: a simulcast terrestrial broadcast mode, and a high data rate cable mode. The two modes share the same pilot, symbol rate, data frame structure, interleaving, Reed-Solomon coding, and synchronization pulses. The terrestrial broadcast mode is optimized for maximum service area, and supports one ATV signal in a 6 MHz channel. The high data rate cable mode, which trades off some robustness for twice the data rate, supports two ATV signals in one 6 MHz channel.

Both modes of the VSB transmission subsystem take advantage of a pilot, a segment sync, and a training sequence for robust acquisition and operation. The two system modes also share identical carrier, sync, and clock recovery circuits, as well as phase correctors and equalizers. Additionally, both modes use the same Reed-Solomon (RS) code for forward error correction (FEC).

In order to maximize service area, the terrestrial broadcast mode incorporates both an NTSC rejection filter (in the receiver) and trellis coding. Pre-coding at the transmitter is incorporated in the trellis code. When the NTSC rejection filter is activated in the receiver, the trellis decoder is switched to a trellis code corresponding to the encoder trellis code concatenated with the filter.

The cable mode, on the other hand, does not have as severe an environment to work in as that of the terrestrial system. Therefore, a higher data rate is transmitted in the form of more data levels (bits/symbol). No trellis coding or NTSC interference rejection filters are employed.

VSB transmission inherently requires only processing the in-phase (I) channel signal, sampled at the symbol rate, thus optimizing the receiver for low cost implementation. The decoder only requires one A/D converter and a real (not complex) equalizer operating at the symbol rate of 10.76 Msamples/s.

The parameters for the two VSB transmission modes are shown in Table 9.1.

9.2 Derivation of the bit rate delivered to a transport decoder by the transmission subsystem

The exact symbol rate of the transmission subsystem is given by:

(1) $4.5/286 \times 684 = 10.76...$ MHz

The symbol rate must be locked in frequency to the transport rate. The transmission subsystem carries 2 information bits per trellis-coded symbol, so the gross payload is:

(2) $10.76... \times 2 = 21.52...$ Mbps

To find the net payload delivered to a decoder it is necessary to adjust (2) for the overhead of the Data Segment Sync, Data Field Sync, and Reed-Solomon FEC.

Upon doing this the net payload bit rate of the 8 VSB terrestrial transmission subsystem becomes:

(3) 21.52... Mbps x 312/313 x 828/832 x 187/207 = 19.28... Mbps

The factor of 312/313 accounts for the Data Field Sync overhead of one Data Segment per field. The factor of 828/832 accounts for the Data Segment Sync overhead of four symbol intervals per Data Segment, and the factor of 187/207 accounts for the Reed-Solomon FEC overhead of 20 bytes per Data Segment.

The calculation of the net payload bit rate of the high data rate cable mode is identical except that 16 VSB carries 4 information bits per symbol. Therefore, the net bit rate is twice that of the 8 VSB terrestrial mode:

(4) 19.28... Mbps x 2 = 38.57... Mbps

To get the net bit rate seen by a transport decoder, however, it is necessary to account for the fact that the MPEG sync bytes are removed from the data stream input to the 8 VSB transmitter. This amounts to the removal of one byte per data segment. These MPEG sync bytes are then reconstituted at the output of the 8 VSB receiver. The net bit rate seen by the transport decoder is:

(5) 19.28... Mbps x 188/187 = 19.39... Mbps

The net bit rate seen by the transport decoder for the high data rate cable mode is:

(6) 19.39... Mbps x 2 = 38.78... Mbps

Table 9.1 Parameters for VSB Transmission Modes

Parameter	Terrestrial Mode	High Data Rate Cable Mode
Channel bandwidth	6 MHz	6 MHz
Excess bandwidth	11.5%	11.5%
Symbol rate	10.76 Msymbols/s	10.76 Msymbols/s
Bits per symbol	3	4
Trellis FEC	2/3 rate	None
Reed-Solomon FEC	T=10 (207,187)	T=10 (207,187)
Segment length	832 symbols	832 symbols
Segment sync	4 symbols per segment	4 symbols per segment
Frame sync	1 per 313 segments	1 per 313 segments
Payload data rate	19.28 Mbps	38.57 Mbps
NTSC co-channel rejection	NTSC rejection filter in receiver	N/A
Pilot power contribution	0.3 dB	0.3 dB
C/N threshold	14.9 dB	28.3 dB

9.3 Performance characteristics of terrestrial broadcast mode

The terrestrial VSB system can operate in a signal-to-additive-white-Gaussian-noise (S/N) environment of 14.9 dB. The 8 VSB, 4-state segment error probability curve in Figure 9.1 shows a segment error probability of 1.93 x 10^{-4}. This is equivalent to 2.5

segment errors/second which has been established by measurement as the threshold of visibility of errors.

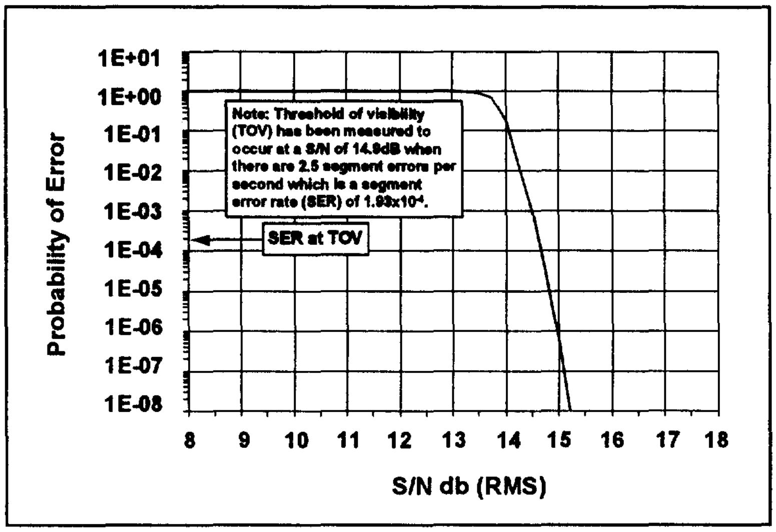

Figure 9.1. Segment error probability, 8 VSB with 4 state trellis, RS (207,187).

The cumulative distribution function (CDF) of the peak-to-average power ratio, as measured on a low power transmitted signal with no non-linearities, is plotted in Figure 9.2. The plot shows that 99.9% of the time the transient peak power is within 6.3 dB of the average power.

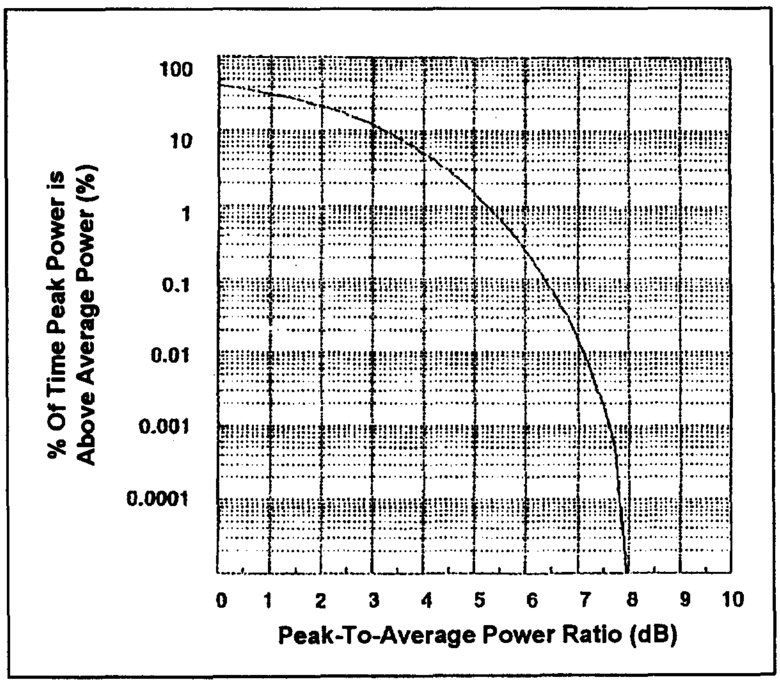

Figure 9.2. Cumulative distribution function of 8 VSB peak-to-average power ratio.

9.4 Transmitter signal processing

A pre-equalizer filter is recommended for use in over-the-air broadcasts where the high power transmitter may have significant in-band ripple or roll off at band edges. This linear distortion can be detected by an equalizer in a reference demodulator ("ideal" receiver) located at the transmitter site that is receiving a small sample of the antenna signal feed provided by a directional coupler which is recommended to be located at the sending end of the antenna feed transmission line. The reference demodulator equalizer tap weights can be transferred into the transmitter pre-equalizer for pre-correction of transmitter linear distortion.

A suitable pre-equalizer is an 80 tap, feed-forward transversal filter. The taps are symbol spaced (93 nsec) with the main tap being approximately at the center, giving approximately ±3.7 μs correction range. It operates on the I channel data signal (there is no Q channel data in the transmitter), and shapes the frequency spectrum of the IF signal so that there is a flat in-band spectrum at the output of the high power transmitter that feeds the antenna for transmission. There is no effect on the out-of-band spectrum of the transmitted signal.

The transmitter VSB filtering may be implemented by complex-filtering the baseband data signal, creating precision-filtered and stable in-phase and quadrature-phase modulation signals. This filtering process provides the root raised cosine Nyquist filtering as well as the sin x/x compensation for the D/A converters. The orthogonal baseband signals are converted to analog form (D/A converters) and then modulated on quadrature IF carriers to create the vestigial sideband IF signal by sideband cancellation (phasing method). The nominal frequency of the IF carrier (and small in-phase pilot) in the prototype hardware used in ACATS testing is 46.69 MHz, which is equal to the IF center frequency (44.000 MHz) plus the symbol rate divided by 4 (10.762 MHz / 4 = 2.6905 MHz). Additional adjacent channel suppression (beyond that achieved by sideband cancellation) may be performed by a linear phase, flat amplitude response SAW filter. Other implementations for VSB filtering are possible which may include the prefilter of the previous section.

9.5 Upconverter and RF carrier frequency offsets

Modern NTSC TV transmitters use a two-step modulation process. The first step usually is modulation of the data onto an IF carrier, which is the same frequency for all channels, followed by translation to the desired RF channel. The VSB transmitter applies this same two-step modulation process. The RF upconverter translates the filtered flat IF data signal spectrum to the desired RF channel. For the same approximate coverage as an NTSC transmitter (at the same frequency), the average power of the ATV signal is 12 dB less than the NTSC peak sync power.

The frequency of the RF upconverter oscillator in ATV terrestrial broadcasts will typically be the same as that used for NTSC (except for NTSC offsets). However, in extreme co-channel situations, the ATV system is designed to take advantage of precise RF carrier frequency offsets with respect to the NTSC co-channel carrier. As the VSB data signal sends repetitive synchronizing information (segment syncs), precise offset causes NTSC co-channel carrier interference into the VSB receiver to phase alternate from sync to sync. The VSB receiver circuits average successive syncs to cancel the interference and make data segment sync detection more reliable.

For ATV co-channel interference into NTSC, the interference is noise-like and does not change with precise offset. Even the ATV pilot interference into NTSC does not benefit from precise frequency offset because it is so small (11.3 dB below the data power) and falls far down the Nyquist slope (20 dB or more) of NTSC receivers.

The ATV co-channel pilot should be offset in the RF upconverter from the dominant NTSC picture carrier by an odd multiple of half the Data Segment rate. A consequential spectrum shift of the VSB signal into the upper adjacent channel is required. An additional offset of 0, +10 kHz, or -10 kHz is required to track the principal NTSC interferer.

For ATV-into-ATV co-channel interference, precise carrier offset prevents possible misconvergence of the adaptive equalizer. If perchance the two ATV Data Field Sync signals should fall within the same data segment time, the adaptive equalizer could misinterpret the interference as a ghost. To prevent this, a carrier offset of $f_{seg}/2$ = 6.47 kHz is recommended for close ATV-into-ATV co-channel situations. This causes the interference to have no effect in the adaptive equalizer.

9.6 Performance characteristics of high speed cable mode

The high data rate cable mode can operate in a signal-to-white-noise environment of 28.3 dB. The error probability curve is shown in Figure 9.3.

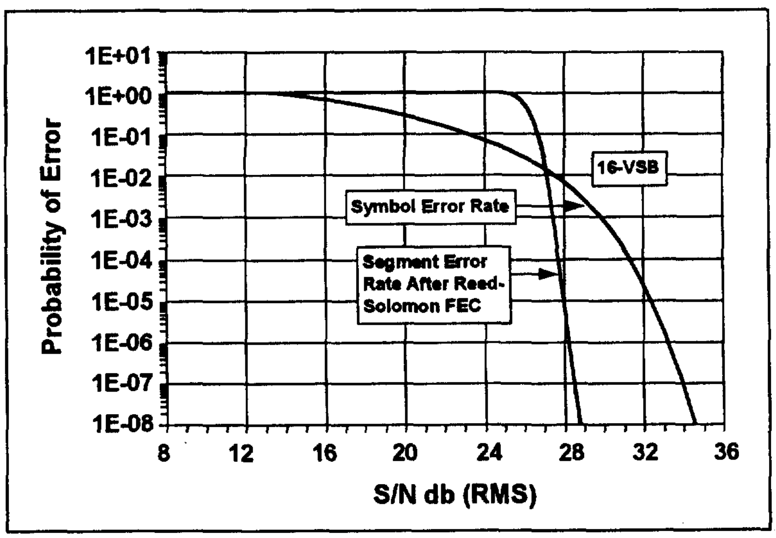

Figure 9.3. 16 VSB error probability.

The cumulative distribution function (CDF) of the peak-to-average power ratio, as measured on a low power transmitted signal with no non-linearities, is plotted in Figure 9.4 and is slightly higher than that of the terrestrial mode.

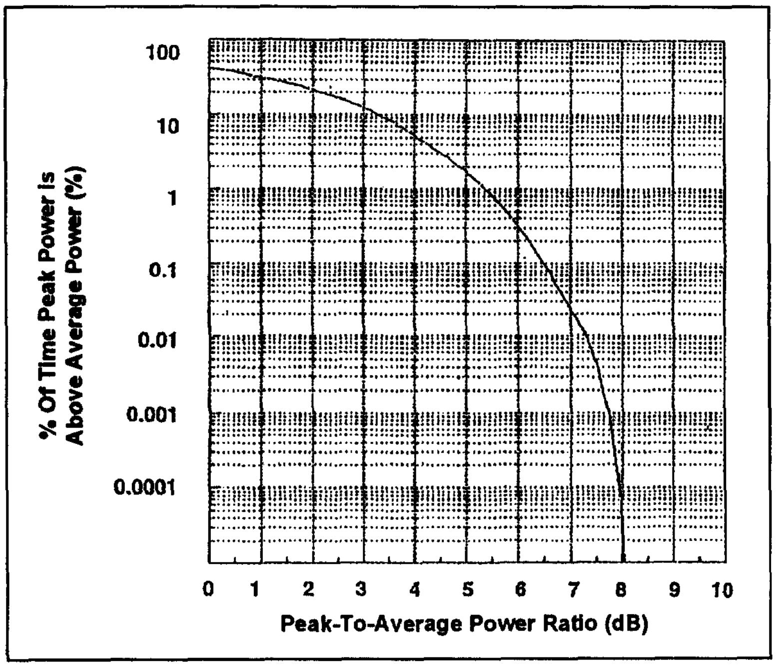

Figure 9.4. Cumulative distribution function of 16 VSB peak-to-average power ratio.

10. RECEIVER CHARACTERISTICS

10.1 Introduction

This Chapter describes the receiver characteristics and design considerations for receivers for the digital television system. It is designed as a tutorial intended for receiver manufacturers to serve as a guideline and literature reference. It illustrates design choices for receivers based on the system implementation example given by the Grand Alliance.

There is no practical implementation of a digital television receiver in existence at the time of release of this Guide. Only a first generation hardware prototype was implemented for laboratory and field tests. The information contained herein will, therefore, not yield a design for a production-ready receiver without substantial development effort by each manufacturer.

In the design of early digital television receivers it will be important to study the test methods, test conditions, and results that were designed and recorded by the Working Parties of the Planning and Systems Subcommittees of ACATS. They attempted to include in the testing all known and theoretically important reception conditions for signals. The data they gathered will allow consumer product designers to prepare a receiver design suitable for most commonly found conditions. Experience gained during the introduction of this new service is likely to teach the industry further important improvements for delivering the service and for receiver design.

10.2 Receiver RF issues

10.2.1 RF characteristics

10.2.1.1 Planning factors used by ACATS PS/WP3

The transmission subsystem is described in the ACATS report of February 24, 1994 under the heading RF/Transmission Characteristics. This summary provides background as to how it was selected and what planning factors were assigned to it.

The selection of a transmission subsystem comprised the following: first, a series of laboratory tests were executed at the Advanced Television Test Center (ATTC) to determine the performance limits of candidate transmission subsystems with respect to channel impairments including noise, interference and multipath. Second, the results of these tests and subsequent VSB modem tests, together with a set of receiver planning factors, were included in a nationwide spectrum utilization computer model which was developed under the direction of the Spectrum Utilization and Alternatives Working Party (PS/WP3) of the Planning Subcommittee of ACATS.

The results of the ATTC tests of transmission subsystems are summarized in Figure 1 of PS/WP3 Document 296 dated February 17, 1994. That figure summarizes the performance of 8 VSB and 32 QAM in the presence of: thermal noise; co-channel and adjacent-channel interference from ATV and NTSC; performance as a taboo interferer into NTSC; and change in noise threshold due to specific ensemble multipath characteristics (which did not enter into the computer model).

The receiver planning factors that entered into the spectrum utilization computer model can be found in the Final Report of PS/WP3 under the heading "Receiver Planning Factors Applicable to All ATV Systems." Table 10.1 was taken from the latest document of PS/WP3, dated December 1, 1994, and shows these Planning Factors. It is also footnoted that "antenna factors are based on the geometric mean frequencies of the three broadcast bands," and that, in addition to F/B, "a formula is employed for the forward lobe simulating an actual receiving antenna pattern."

Table 10.1 Receiver Planning Factors Used by PS/WP3

Planning factors	Low VHF	High VHF	UHF
Antenna impedance (ohms)	75	75	75
Bandwidth (MHz)	6	6	6
Thermal noise (dBm)	-106.2	-106.2	-106.2
Noise figure: (dB)	10	10	10
Frequency (MHz)	69	194	615
Antenna factor (dBm/dBµ)	-111.7	-120.7	-130.7
Line loss (dB)	1	2	4
Antenna gain (dB)	4	6	10
Antenna F/B ratio (dB)	10	12	14

Taking both the ATTC results and receiver planning factors into account, the computer analysis was done which led PS/WP3 to conclude, in its Document 296, that, "as the natural outcome of employing as planning factors the results of the ATTC interference tests, the 8 VSB transmission subsystem shows an advantage over the 32 QAM subsystem both during the transition period when the spectrum is shared with NTSC and after the transition when only ATV will be broadcast."

The computer model calculates signal conditions at locations throughout the country, based upon transmitter locations, power, propagation models, etc. The performance numbers—such as C/N or D/U ratios—that entered into the computer model as the result of the ATTC tests will, in the real world, be a function of these signal conditions and the entire receiver installation including actual antenna, feedline and receiver performance.

10.2.1.2 Noise figure

A number of factors enter into the ultimate carrier-to-noise ratio within the receiver. For example, the receiver planning factors applicable to UHF used by PS/WP3 (reference Final Report of the Spectrum Utilization and Alternatives Working Party of the Planning Subcommittee of the Advisory Committee on Advanced Television Service) are shown in Table 10.1.

A consumer can effect an improvement in the noise performance of his installation by impacting any contributing factor; examples include installation of a low-noise amplifier (LNA) or a better antenna.

10.2.1.3 Co-channel and adjacent-channel rejection

The assumptions made by PS/WP3 as reported in its Document 296 were based on threshold-of-visibility (TOV) measurements determined during the test of the competing systems. These TOV numbers were correlated to BER test results from the same tests. When testing the Grand Alliance VSB modem hardware at ATTC, and in field tests, only BER measurements were taken. In Table 10.2 the results of these tests are again expressed in equivalent TOV numbers derived from the BER measurements

Table 10.2 ATV Interference Criteria

Co-channel ATV-into-NTSC	33.8 dB
Co-channel NTSC-into-ATV	2.07 dB
Co-channel ATV-into-ATV	15.91 dB
Upper-adjacent ATV-into-NTSC	-16.17 dB
Upper-adjacent NTSC-into-ATV	-47.05 dB
Upper-adjacent ATV-into-ATV	-42.86 dB
Lower-adjacent ATV-into-NTSC	-17.95 dB
Lower-adjacent NTSC-into-ATV	-48.09 dB
Lower-adjacent ATV-into-ATV	-42.16 dB

It should be noted that the exact amount of adjacent-channel or co-channel interference entering the receiver terminals is a function of the exact overall antenna gain pattern used (not just F/B ratio) which is also a function of frequency.

10.2.1.4 Unintentional radiation

This subject is already well covered by the FCC Rules, Part 15.

10.2.1.5 Direct pickup (DPU)

Rules for direct pickup are included in the FCC Rules, Part 15. As advanced television will use digital transmission, rather than analog transmission, the image rejection numbers of IS-23 Section 3.28 of 50 dB or 60 dB may be far too stringent for the digital broadcasting case. A digital system is much more robust against image interference, and regulation of image rejection depth will not be necessary. Again the assumptions of IS-23 based on an IF of 45 MHz must be re-examined.

10.2.2 The modem field test

A preliminary version of the Grand Alliance transmission subsystem was provided for field testing in the summer of 1994. Terrestrial testing was performed using the 8 VSB mode and cable testing was performed using the 16 VSB high data rate cable mode.

Results of the tests are included in a September 16, 1994 ACATS report "*Field Test Results of the Grand Alliance Transmission Subsystem,*" Document SS/WP2-1354.

The test results contain valuable information for receiver manufacturers about multipath interference and other impairment conditions, and their effect on the bit error rate of the digital signal.

10.2.3 Signal conversion and carrier recovery.

The following descriptions were taken from the Grand Alliance HDTV System Specification and are specific to the hardware implementation of the Grand Alliance. Figure 10.1 shows the receiver block diagram of the VSB terrestrial broadcast transmission system. Descriptions of each block follow.

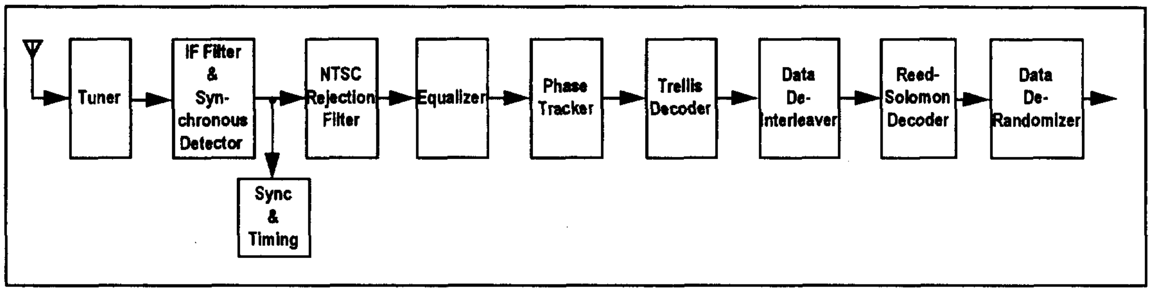

Figure 10.1. VSB receiver.

10.2.3.1 Tuner

The tuner, illustrated in Figure 10.2, as implemented in the prototype submitted for test, receives the 6 MHz signal (UHF or VHF) from the antenna. It is a high-side injection double-conversion type with a first IF frequency of 920 MHz. This puts the image frequencies above 1 GHz, making them easy to reject by a fixed front end filter. This selection of first IF frequency is high enough so that the input band-pass filter selectivity prevents the local oscillator (978-1723 MHz) from leaking out the tuner front end and interfering with other UHF channels, yet it is low enough for second harmonics of UHF channels (470-806 MHz) to fall above the first IF band-pass. Harmonics of cable channels could possibly occur in the first IF pass-band but are not a real problem because of the relatively flat spectrum (within 10 dB) and small signal levels (-28 dBm or less) used in cable systems.

The tuner input has a band-pass filter that limits the frequency range to 50-810 MHz, rejecting all other non-television signals that may fall within the tuner's image frequency range (beyond 920 MHz). In addition, a broadband tracking filter rejects other television signals, especially those much larger in signal power than the desired signal power. This tracking filter is not narrow, nor is it critically tuned, as is the case of present day NTSC tuners that must reject image signals only 90 MHz away from the desired channel. Minimal channel tilt, if any, exists due to this tracking filter.

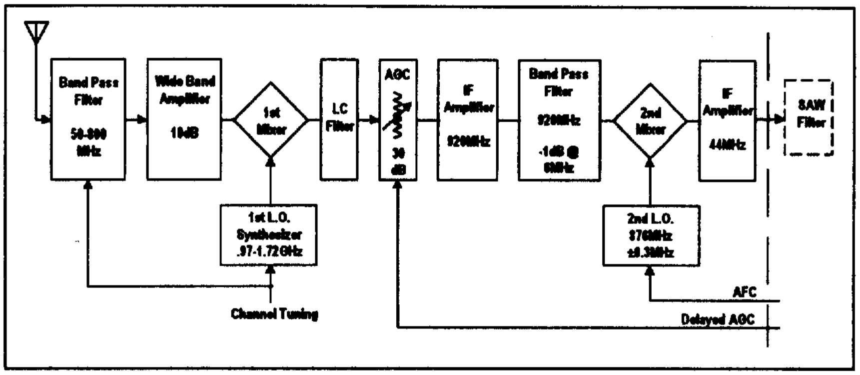

Figure 10.2. Tuner block diagram.

A 10 dB gain, wideband RF amplifier increases the signal level into the first mixer, and is the dominant determining factor of receiver noise figure (7-9 dB over entire VHF, UHF, and cable bands). The first mixer, a highly linear double-balanced design to minimize even harmonic generation, is driven by a synthesized low phase noise local oscillator (LO) above the first IF frequency (high-side injection). Both the channel tuning (first LO) and broadband tracking filtering (input band-pass filter) are controlled by microprocessor. The tuner is capable of tuning the entire VHF and UHF broadcast bands as well as all standard, IRC, and HRC cable bands.

The mixer is followed by an LC filter in tandem with a narrow 920 MHz band-pass ceramic resonator filter. The LC filter provides selectivity against the harmonic and sub-harmonic spurious responses of the ceramic resonators. The 920 MHz ceramic resonator band-pass filter has a -1 dB bandwidth of about 6 MHz. A 920 MHz IF amplifier is placed between the two filters. Delayed AGC of the first IF signal is applied immediately following the first LC filter. The 30 dB range AGC circuit protects the remaining active stages from large signal overload.

The second mixer is driven by the second LO, which is an 876 MHz voltage-controlled SAW oscillator. It is controlled by the frequency and phase-locked loop (FPLL) synchronous detector. The second mixer, whose output is the desired 44 MHz second IF frequency, drives a constant gain 44 MHz amplifier. The output of the tuner feeds the IF SAW filter and synchronous detection circuitry.

The tuner is made out of standard consumer electronic components, and is housed in a stamped metal enclosure.

10.2.3.2 Channel filtering and VSB carrier recovery

Carrier recovery is performed on the small pilot carrier by an FPLL circuit, illustrated in Figure 10.3. The first LO is synthesized by a PLL and controlled by a microprocessor. The third LO is a fixed reference oscillator. Any frequency drift or deviation from nominal has to be compensated in the second LO. Control for the second LO comes from the FPLL synchronous detector, which integrally contains both a frequency loop and a phase-locked loop in one circuit. The frequency loop provides a wide frequency pull-in range of ±100 kHz while the phase-locked loop has a narrow bandwidth (less than 2 kHz).

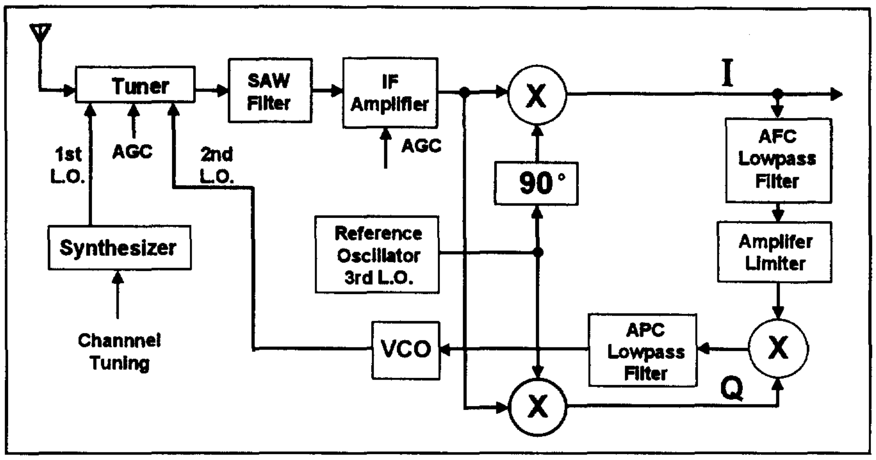

Figure 10.3. Tuner - IF - FPLL.

During frequency acquisition, the frequency loop uses both the in-phase (I) and quadrature-phase (Q) pilot signals. All other data processing circuits in the receiver use only the I-channel signal. Prior to phase-lock, as is the condition after a channel change, the automatic frequency control (AFC) low-pass filter acts on the beat signal created by the frequency difference between the VCO and the incoming pilot. The high frequency data (as well as noise and interference) is mostly rejected by the AFC filter, leaving only the pilot beat frequency. After limiting this pilot beat signal to a constant amplitude (±1) square wave, and using it to multiply the quadrature signal, a traditional bipolar S-curve AFC characteristic is obtained. The polarity of the S-curve error signal depends upon whether the VCO frequency is above or below the incoming IF signal. Filtered and integrated by the automatic phase control (APC) low-pass filter, this DC signal adjusts the tuner's second LO to reduce the frequency difference.

When the frequency difference comes close to zero, the APC loop takes over and phase-locks the incoming IF signal to the third LO. This is a normal phase-locked loop circuit, with the exception that it is bi-phase stable. However, the correct phase-lock polarity is determined by forcing the polarity of the pilot to be equal to the known transmitted positive polarity. Once locked, the detected pilot signal is constant, the limiter output feeding the third multiplier is at a constant +1, and only the phase-locked loop is active (frequency loop automatically disabled). The APC low-pass filter is wide enough to reliably allow ±100 kHz frequency pull-in, yet narrow enough to consistently reject all strong white noise (including data) and NTSC co-channel interference signals. The PLL has a bandwidth that is narrow enough to reject most of the AM and PM generated by the data, yet is wide enough to track out any phase noise on the signal (and, hence, on the pilot) out to about 2 kHz. Tracking out low frequency phase noise (as well as low frequency FM components) allows the phase tracking loop, discussed in Section 10.2.3.8, to be more effective.

The prototype receiver can acquire a signal and maintain lock at a signal-to-noise ratio of 0 dB or less, and in the presence of heavy interference.

10.2.3.3 Segment sync and symbol clock recovery

The repetitive data segment syncs (Figure 10.4) are detected from among the synchronously detected random data by a narrow bandwidth filter. From the data segment syncs, a properly phased 10.76 MHz symbol clock is created along with a coherent AGC control signal. A block diagram of this circuit is shown in Figure 10.5.

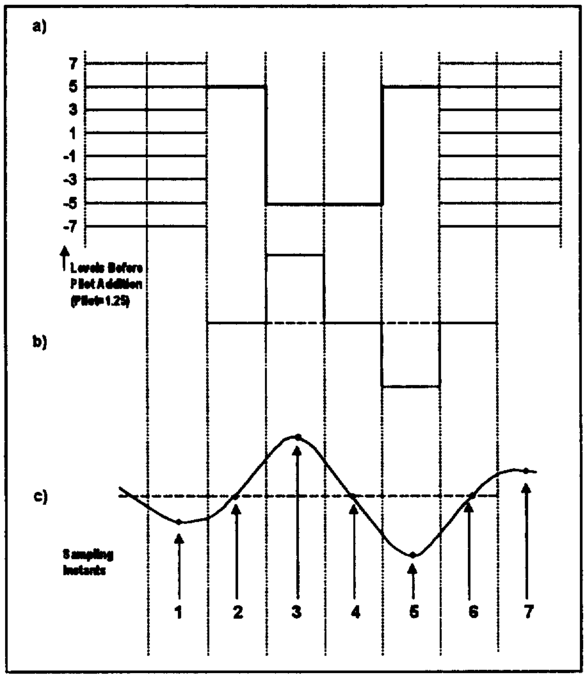

Figure 10.4. Data segment sync.

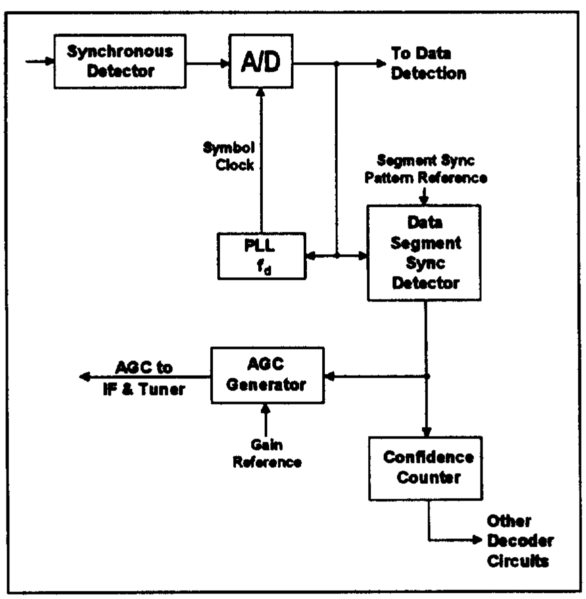

Figure 10.5. Segment sync & symbol clock recovery with AGC.

The 10.76 Msymbols/s (684 ÷ 286 x 4,500,000 Hz) I-channel composite baseband data signal (syncs and data) from the synchronous detector is converted by an A/D converter for digital processing. Traditional analog data eyes can be viewed after synchronous detection. However, after conversion to a digital signal, the data eyes cannot be seen due to the sampling process. A PLL is used to derive a clean 10.76 MHz symbol clock for the receiver.

With the PLL free-running, the data segment sync detector containing a 4-symbol sync correlator looks for the two level syncs occurring at the specified repetition rate. The repetitive segment sync is detected while the random data is not, enabling the PLL to lock on the sampled sync from the A/D converter, and achieve data symbol clock synchronization. Upon reaching a predefined level of confidence (using a confidence counter) that the segment sync has been found, subsequent receiver loops are enabled.

Data segment sync detection and clock recovery both work reliably at signal-to-noise ratios of 0 dB or less, and in the presence of heavy interference.

10.2.3.4 Non-coherent and coherent AGC

Prior to carrier and clock synchronization, non-coherent automatic gain control (AGC) is performed whenever any signal (locked or unlocked signal, or noise/interference) overruns the A/D converter. The IF and RF gains are reduced accordingly, with the appropriate AGC "delay" applied.

When data segment syncs are detected, coherent AGC occurs using the measured segment sync amplitudes. The amplitude of the bipolar syncs, relative to the discrete levels of the random data, is determined in the transmitter. Once the syncs are detected in the receiver, they are compared to a reference value, with the difference (error) integrated. The integrator output then controls the IF and "delayed" RF gains, forcing them to whatever values provide the correct sync amplitudes.

10.2.3.5 Data field synchronization

Data Field Sync detection, shown in Figure 10.6, is achieved by comparing each received data segment from the A/D converter (after interference rejection filtering to minimize co-channel interference) with ideal field #1 and field #2 reference signals in the receiver. Over-sampling of the field sync is NOT necessary as a precision data segment and symbol clock has already been reliably created by the clock recovery circuit. Therefore, the field sync recovery circuit knows exactly where a valid field sync correlation should occur within each data segment, and only needs to perform a symbol by symbol difference. Upon reaching a predetermined level of confidence (using a confidence counter) that field syncs have been detected on given data segments, the Data Field Sync signal becomes available for use by subsequent circuits. The polarity of the middle of the three alternating 63 bit pseudo random (PN) sequences determines whether field 1 or field 2 is detected.

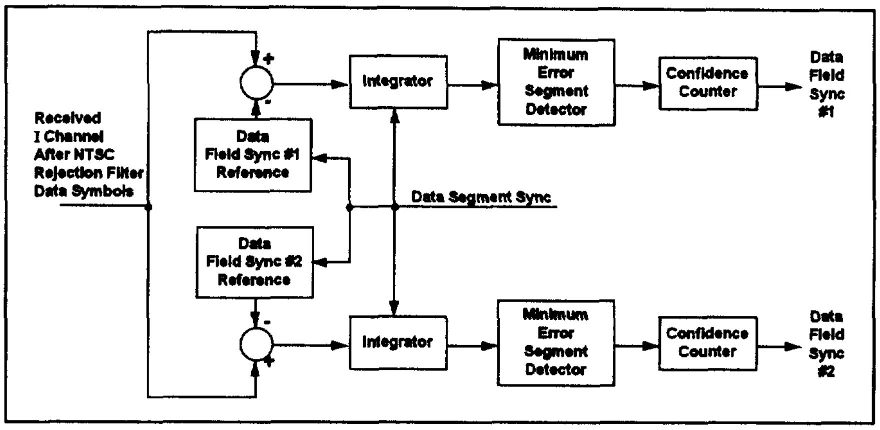

Figure 10.6. Data field sync recovery.

This procedure makes field sync detection robust, even in heavy noise, interference, or ghost conditions. Field sync recovery can reliably occur at signal-to-noise ratios of 0 dB or less, and in the presence of heavy interference.

10.2.3.6 Interference rejection filter

The interference rejection properties of the VSB transmission system are based on the frequency location of the principal components of the NTSC co-channel interfering signal within the 6 MHz television channel and the periodic nulls of a VSB receiver baseband comb filter.

Figure 10.7a shows the location and approximate magnitude of the three principal NTSC components: (1) the visual carrier (V) located 1.25 MHz from the lower band edge, (2) the chrominance subcarrier (C) located 3.58 MHz higher than the visual carrier frequency, and (3) the aural carrier (A) located 4.5 MHz higher than the visual carrier frequency.

The NTSC interference rejection filter (comb) is a one tap linear feed-forward filter, as shown in Figure 10.8. Figure 10.7b shows the frequency response of the comb filter, which provides periodic spectral nulls spaced 57 * f_H (10.762 MHz / 12, or 896.85 kHz) apart. There are 7 nulls within the 6 MHz channel. The NTSC visual carrier frequency falls close to the second null from the lower band edge. The 6th null from the lower band edge is correctly placed for the NTSC chrominance subcarrier, and the 7th null from the lower band edge is near the NTSC aural carrier.

Comparing Figure 10.7a and Figure 10.7b shows that the visual carrier falls 2.1 kHz below the second comb filter null, the chroma subcarrier falls near the 6th null, and the aural carrier falls 13.6 kHz above the 7th null. (Note, the aural carrier is at least 7 dB below its visual carrier).

The comb filter, while providing rejection of steady-state signals located at the null frequencies, has a finite response time of 12 symbols (1.115 μs). Thus, if the NTSC interfering signal has a sudden step in carrier level (low to high or high to low), one cycle of the zero-beat frequency (offset) between the ATV and NTSC carrier frequencies will pass through the comb filter at an amplitude proportional to the NTSC step size as instantaneous interference. Examples of such steps of NTSC carrier are the leading and trailing edge of sync (40 IRE units). If the desired to undesired (D/U) signal power ratio is large enough, data slicing errors will occur. However, interleaving will spread the interference and will make it easier for the Reed-Solomon code to correct them (RS can correct up to 10 byte errors/segment).

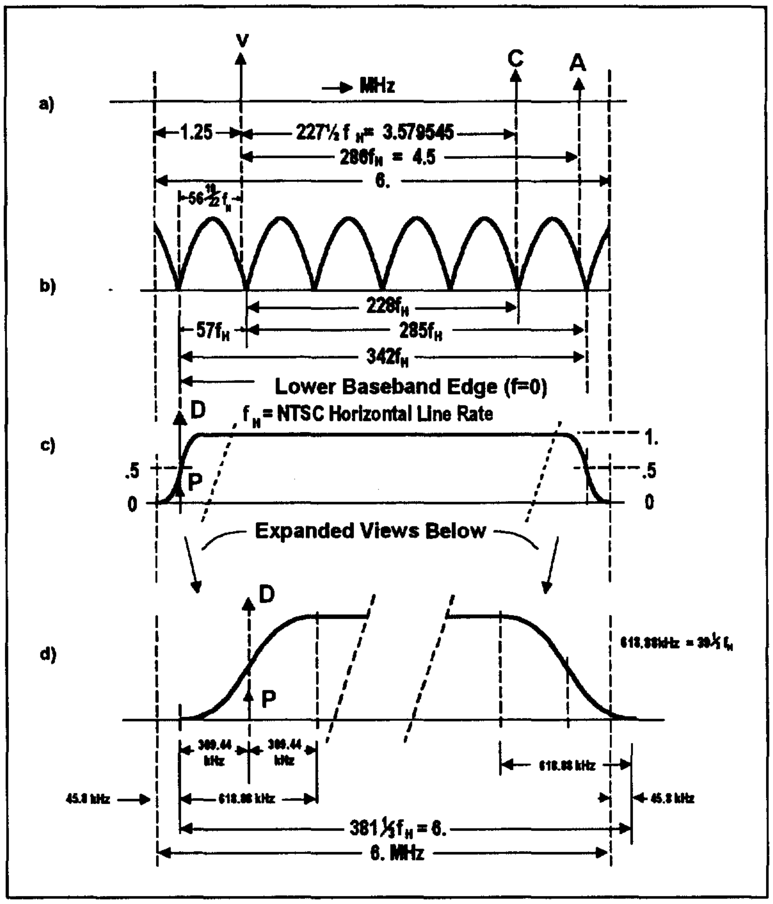

Figure 10.7. Location of NTSC carriers - comb filtering.

Although the comb filter reduces the NTSC interference, the data is also modified. The 7 data eyes (8 levels) are converted to 14 data eyes (15 levels). This conversion is caused by the partial response process which is a special case of intersymbol interference that does not close the data eye but creates double the number of eyes of the same magnitude. The modified data signal can be properly decoded by the trellis decoder, and will be described in Section 10.2.3.9. Note that, because of time sampling, only the maximum data eye value is seen after A/D conversion.

The detail at the band edges for the overall channel is shown in Figure 10.7c and Figure 10.7d. Figure 10.7d shows that the frequency relationship of 56 $^{19}/_{22}$ * f_H between the NTSC visual carrier and the ATV carrier requires a shift in the ATV spectrum with respect to the nominal channel. The shift equals +45.8 kHz, or about +0.76%. This is slightly higher than currently applied channel offsets and reaches into the upper adjacent-channel at a level of about -40 dB. If that is another ATV channel, its spectrum is also